Working, Consuming, and Dying: Quantifying the Diversity in the American Experience

Chadwick Curtis*
University of Richmond

Julio Garín†
Claremont McKenna College

Robert Lester‡
Colby College

This Version: March 7, 2022

Abstract

We document how lifetime utility varies by demographic groups in the US and how these differences have evolved since the start of the 21st century. Using the equivalent variation as our measure of welfare we find that the standard deviation in cross-sectional well-being between demographic groups is comparable to the standard deviation of relative annual income in prime earning years and double the standard deviation of relative consumption. Our metric includes consumption, leisure, and mortality risk. The results are primarily driven by differences in consumption and life expectancy. Controlling for other demographics, welfare is increasing in educational attainment and is higher for women and those of Asian descent. This qualitative ordering is robust to classifying a broad measure of home production and child care as work and various definitions of real consumption. Finally, we show that changes in mortality rates associated with 'deaths of despair' disproportionately lower the welfare of less educated Whites.

JEL Classification: I31; 051.

Keywords: Welfare; Life-cycle; Inequality; Demographics.

---

*Corresponding author at: Deapartment of Economics, Robins School of Business, University of Richmond, 1 UR Dr., Richmond, VA 23173, USA. *E-mail address:* ccurtis2@richmond.edu.

†*E-mail address:* jgarin@gmail.com.

‡*E-mail address:* rblester@colby.edu.

1 Introduction

The facts of cross-sectional inequality across a number of dimensions are well established. Consumption, leisure, and life expectancy vary substantially between and within narrowly defined demographic groups. In this paper we aim to answer the following question: Accounting for these well-developed metrics of well-being, how much does the welfare of individuals belonging to various racial, educational, and gender groups compare to the welfare of the average American?

When looking at the data, it is not always clear how the measurable inputs to well-being translate to overall welfare. For instance, those groups with high consumption generally have lower leisure. Moreover, despite being higher within genders, the overall correlation between life expectancy and consumption is less than 0.5. Thus, inferring economic well-being from the underlying inputs requires a way of ranking competing bundles of consumption, leisure, and life expectancy, i.e. a utility function. It is the combination of data with a lifetime utility function that allows us to make welfare comparisons across groups and over time.

Following Jones and Klenow (2016) (JK henceforth), we use the equivalent variation – the amount of consumption you must compensate or tax the average American to be indifferent between living life as themselves or an average member of a specific demographic group – as our measure of welfare. Whereas JK study economic welfare across countries, we focus on the distribution of welfare within the United States. A person's utility within a period depends on consumption and leisure. Each person also receives a set of group-specific mortality rates over their life cycle which gives rise to different life expectancies across demographic groups. Thus, welfare dispersion is driven by differences in average consumption, average leisure, and mortality rates.

We uncover a number of facts. First, we find that welfare is higher for women than men, highest for Asians and lowest for Blacks, and increasing in educational attainment. While men spend more time on market work than women on average, we show that the welfare "gender gap" persists even when classifying child care and home production as non-leisure activities. Second, the standard deviation of economic welfare is comparable to the standard deviation of income at midlife and more than double the standard deviation of consumption. There is also a wide range, with college-educated Asian women having more than three times the welfare of Non-Hispanic White and Black men without a college degree. Third, differences in life expectancy and consumption drive the welfare results with differences in leisure playing a more minor role. Fourth, the correlation between welfare and income is about 0.64, but rises to more than 0.9 within each gender. Finally, economic welfare for the average American grew by about 30 percent cumulatively, or two percent annually, between the first five years of this century and 2013-2017. Welfare growth is especially high for Black men and for individuals with a college degree. There is a correlation of about -0.26 between initial welfare levels and subsequent growth, implying a limited degree of welfare convergence.

Although average welfare has increased since the start of the century, welfare growth due to changes in life expectancy is the smallest for less than college-educated White men and women without a college degree. To further study this finding, and motivated by the increase in "deaths of despair" since the turn of the century documented by Case and Deaton (2017), we calculate the welfare gains (and the distribution of those gains) from eliminating the rise in deaths caused by suicide, drug overdose, liver failure or cirrhosis since 2000-2004. We find that Whites without college degrees would be willing to sacrifice the most consumption to eliminate deaths of despair. Less educated White men, the hardest hit group, would be willing to give up about seven percent of lifetime consumption to return to the early 2000s deaths of despair mortality rates. This finding is consistent with Case and Deaton (2017) who find that the increasing prevalence of deaths of despair is enough to decrease the life expectancies of non-Hispanic Whites without a college education, while college-educated Whites and members of other groups, irrespective of education, have continued to see improvements in life expectancy.

Our paper relates to a long-standing literature attempting to create plausible proxies of economic welfare and then using those proxies to understand the distribution of welfare across countries and how it has changed over time within countries. An entire NBER volume in 1973 was devoted to studying how economic and social performance has evolved in the US (Moss, 1973). Anticipating future work in the subject, Nordhaus and Tobin (1971) build a measure of economic welfare that includes consumption, time use, and urban (dis)amenities.[1] More recently, JK build a measure of economic welfare that includes consumption, leisure, life expectancy, and inequality. They show that cross-country welfare is strongly correlated with GDP per capita and that welfare has been growing faster than income per person on average. They note that the high correlation in the cross section hides some discrepancies between welfare and income. Falcettoni and Nygaard (2020) use the JK utility function to study how economic welfare is distributed across states in the US. Like us, they find that there is significant dispersion in welfare. While we share a common framework with Falcettoni and Nygaard (2020), we focus on different units on observation – with ours being demographic groups and theirs being states. Brouillette, Jones, and Klenow (2021) quantify welfare differences between White and Black Americans. Like us, they find that welfare for the average Black American is lower than the average White American, but that this difference is closing over time. Unlike Brouillette, Jones, and Klenow (2021), we disaggregate the results by education and show that the Black-White welfare gap is higher for the college educated.

Our paper is also related to a voluminous literature studying various measures of economic inequality in the US. Trends in income inequality have been extensively studied and are summarized in Goldin and Katz (2008) and Acemoglu and Autor (2011). The general consensus is that there exists a sizable college wage premium that has grown since the late 70s; real wages have fallen slightly over the same time period for workers without a post-secondary education, and that the increase in inequality has been concentrated in the top relative to the median rather than the median relative to the bottom. While cross-

[1]The idea here is that living in a city entails tolerating more congestion. Had Tobin been alive to witness the proliferation of microbreweries and kombucha bars, perhaps he would have been more inclined to emphasize the upside of living in cities.

sectional consumption inequality is unequivocally lower than income inequality, the trend in consumption inequality is more controversial. Krueger and Perri (2006) show that the large increase in wage income inequality in the last quarter of the 20th century was not accompanied by nearly as much of an increase in consumption inequality. This conclusion has been challenged by Aguiar and Bils (2015) among others who argue that the Consumer Expenditure Survey (CE) is subject to substantial measurement error and that once one accounts for measurement error, consumption inequality has risen about as much as income inequality. In our baseline, we use the CE interview survey and consider some alternative definitions of consumption. These adjustments do not much change the quantitative results.

In addition to consumption and income, a number of recent papers document trends in time use and inequality in life expectancy. Aguiar and Hurst (2007, 2009) show that less educated individuals enjoy more leisure than highly educated individuals and this gap has expanded over time. Chetty et al. (2016) show that income and life expectancy are highly correlated and that differences in life expectancy by income group expanded between 2001 and 2014. Case and Deaton (2017) discuss the rising mortality rates for middle age White non-Hispanic men and women between 1999 and 2013, attributing much of the increased mortality rates to "deaths of despair." Novosad et al. (2020) show that these increasing mortality rates are concentrated among the lowest ten percent of the education distribution. Our paper unifies the differences in cross-group consumption, time use, and mortality rates into a comprehensive welfare measure.

Finally, while our paper is about measurement, there is a large literature investigating the causes and consequences of rising income inequality in structural models with Heathcote, Storesletten, and Violante (2010b) being one example. Beyond some speculation, we are entirely silent on the economic mechanisms that give rise to the distribution of welfare. The quantified welfare differences that we present in this paper can be used in structural models that speak to specific mechanisms and evaluate welfare effects of particular policies. In that regard, one of our goals is to present a way of looking at the data that will be useful for structural modeling.

2 Framework

Our exercise considers the lifetime welfare of a person at age 25 who is assigned a race/ethnicity, gender, and education level. We refer to each unique combination of race/ethnicity, gender, and education level as a 'demographic group.' Lifetime utility (or welfare) is from the perspective of the 25 year old facing the mortality rate, mean consumption and leisure at each age in their demographic classification. The measure of welfare comparisons across groups is the annual equivalent variation: At age 25, how much would the consumption of the average American need to be adjusted to equate the lifetime utility of the average American (with the adjusted consumption profile) to the lifetime utility for a member of a particular demographic group.

We choose to start at age 25 because educational investment is usually complete by this age. Consequently, we are not taking a stand on to what extent the cost of college may affect lifetime resources or how the consumption value of college affects lifetime utility. This allows us to treat education as an endowment rather than a choice and to measure lifetime utility taking education as given.

Following JK, all individuals share a common per-period utility function $u = \bar{u} + \ln c + v(l)$ where c is consumption and l is leisure. The function v is increasing and concave in l. $\bar{u}$ is a constant which ensures flow utility is non-negative. Intuitively, if flow utility was negative, an added year of life expectancy would decrease utility. Adding the constant ensures that the value of additional life year is never negative. Our consumption measure is expenditures on market nondurable goods and services. Leisure is defined as the residual of the time endowment minus the sum of market work, home production, and childcare time. Since men spend relatively more time at work than women and women spend relatively more time at home production and childcare than men, gender disparities in leisure are smaller under our definition of leisure compared to one that that excludes home production and childcare from work. We show how this assumption affects our results in the robustness section.

We begin to track individuals at age 25 after education decisions have been completed. Each person belongs to a particular gender, race, and educational attainment category. We index each gender, race, and education level combination by i and age by a. Lifetime utility for a member of group i is

$$\bar{V}_i = \sum_{a=25}^{99} \beta^{a-25} S_i(a) \left[\bar{u} + \ln c_i(a) + g(a-25) + v(l_i(a))\right] \tag{1}$$

where g is the annual growth rate of consumption and $S(a_i)$ are cumulative survival probabilities, i.e. the probability of surviving through age a. $S(a_i)$ depends on age and group and we assume $S(a_i) = 0$, for $a_i \geq 100$, $\forall i$.[2] We define the equivalent variation as the amount one would have to raise or lower the average American's consumption so that the lifetime utility of the average American with the revised consumption profile equals the lifetime utility of a member of group i. Formally, the equivalent variation, λ, solves the equation

$$\sum_{a=25}^{99} \beta^{a-25} S_i(a) \left\{\bar{u} + \ln c_i(a) + g(a-25) + v\left[l_i(a)\right]\right\} =$$
$$\sum_{a=25}^{99} \beta^{a-25} S_{US}(a) \left\{\bar{u} + \ln\left[c_{US}(a)(1+\lambda_i)\right] + g(a-25) + v\left[l_{US}(a)\right]\right\}$$

where the US subscripts denote the value of each variable for the average American.

Solving for $\ln(1+\lambda_i)$ gives

$$\ln(1+\lambda_i) = \frac{\bar{V}_i - \bar{V}_{US}}{\sum_a \beta^{a-25} S_{US}(a)}. \tag{2}$$

$\lambda_i * 100$ is our measure of economic welfare. By definition, $\lambda_{US} = 0$. Because the flow utility function is additively separable, we can decompose economic welfare into three different

[2] Note we are assuming $S(24) = 1 \forall i$. Thus, we any reference to "life expectancy" is life expectancy conditional to living to 24.

components:[3]

$$\ln(1+\lambda_i) = \underbrace{\sum_a \Delta s_i(a) u_i(a)}_{\text{Life expectancy}} + \underbrace{\sum_a s_{US}(a)\left(u(c_i(a)) - u(c_{US}(a))\right)}_{\text{Consumption}} + \underbrace{\sum_a s_{US}(a)\left(v(l_i(a)) - v(l_{US}(a))\right)}_{\text{Leisure}} \tag{3}$$

where $\Delta s_i(a)$ and $s_{US}(a)$ are defined as

$$\Delta s_i(a) = \frac{\beta^{a-25} S_i(a) - \beta^{a-25} S_{US}(a)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)} \tag{4}$$

$$s_{US}(a) = \frac{\beta^{a-25} S_{US}(a)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)} \tag{5}$$

The derivation is contained in Appendix A. In a robustness exercise, we report results using a compensating variation, i.e. how much one would have to change consumption of group X to be indifferent between living the life path of a member of group X with the revised level of consumption or the life path of the average American. The main difference comes in the life expectancy component. Under the equivalent variation, the life expectancy component is weighted by group i's flow utility. If the group has low flow utility, adding additional life years will not add much to welfare. Conversely, if the group has high flow utility, the life expectancy term will matter more. The compensating variation replaces $u_i(a)$ with $u_{US}(a)$. The life expectancy term will then matter more for groups with relatively low flow utilities compared to those groups with relatively high flow utilities.

[3] In a robustness exercise, we decompose welfare to additionally incorporate within-group consumption and leisure inequality.

3 Data and Calibration

3.1 Data

The primary sample is from 2013-2017. Our demographic groups are combinations of gender, racial/ethnic, and education groups. We use two genders (male and female) and two education groups: less than a bachelor's degree and bachelor's degree or more. We refer to the first group as "high school" or "less than college" group and the other as the "college" group.[4] We also define four race/ethnicity categories: White non-Hispanic, Hispanic, Black non-Hispanic, and Asian non-Hispanic. This results in 16 unique gender-race-education groups. Because of sample size issues, we group everyone into non-overlapping five-year age bins beginning at age 25. Hence, every individual in a given age bin receives the same flow utility.

Our data on consumption comes from the Consumer Expenditure Survey (CE). The CE consists of an interview component in which households are interviewed over five consecutive quarters. In the first quarter surveyors collect demographic and economic characteristics of the household and then track detailed consumption patterns in the subsequent four quarters. The CE also consists of a diary survey in which respondents are asked to record their purchases. Bee, Meyer, and Sullivan (2015) show that the diary survey fails to track the Personal Consumption Expenditure data in the National Accounts for most product categories. On the other hand, they show that the interview survey tracks most of the largest consumption categories in the PCE quite well.[5] With this in mind, we use the interview survey for the consumption data. Some groups have a rather low sample size for individual years, so we merge subjects into five-year age bins from ages 25 to 99. The unweighted consumption sample size for each group for 2013-2017 are shown in Table C2 of Appendix C.

Consumption is measured at the household level. To transform into an individual level, we

[4] Throughout the paper, we refer to the lower education category as "less than college" or "high school" even though many of the people in this classification have attended some college or received a two-year degree.

[5] Some of the exceptions include food purchased away from home and furniture.

divide household consumption by OECD equivalent scales. We calculate real consumption by dividing nominal consumption by the good-specific consumer price index.[6] As is well known, the CE under counts consumption relative to the NIPA.[7] To account for this, we adjust consumption by a scaling factor such that average real per capita consumption on nondurables and services in our sample equals the corresponding figure in the NIPA.[8] Because we want to track consumption rather than expenditure, our definition of consumption excludes purchases of durable goods. We do, however, include the imputed rent of owner occupied housing in our consumption measure. In a robustness exercise, we report the results assuming durable good expenditure counts as consumption. Due to age censoring in the CE, consumption expenditures for older ages are grouped into one age category and we cannot separate by five-year age groups past 80. Therefore, for those age 85 and older, we forecast consumption by regressing group specific consumption on a quadratic in age and then using the group-specific coefficients to predict consumption. This method is presented in Appendix C.

We use the American Time Use Survey (ATUS) to calculate our leisure. Non-leisure time is calculated as annual hours worked (h_{work}) plus hours in home production (h_{home}) – producing goods and services at home that could alternatively be purchased in the market – plus hours providing childcare (h_{child}). Leisure is then $1 - \frac{h_{work}+h_{home}+h_{child}}{365\times 16}$. Appendix C contains a detailed description on this calculation and the components of time use.

Figure 1 decomposes non-leisure time across the various demographic groups. Conditional on race and education, men spend more time on market work than women. On the other hand, women generally spend more time doing home production and childcare activities than men. When adding up all non-leisure activities, the figure shows that the differences in labor/leisure time between sexes are relatively small within each race and education group.

[6] See Appendix C for a complete list of the categories.

[7] See Heathcote et al. (2010a) for a discussion of this.

[8] We use the CPI to do this because it is common to the CE and the NIPA.

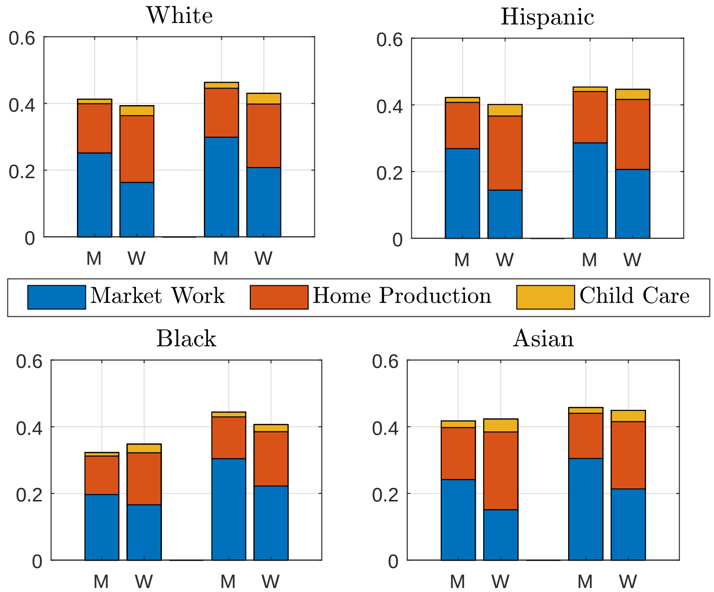

Figure 1: *Market Work and Home Work.*

Notes: Decomposition of share of time spent in market work, home production, and child care. Components of home work and child care are detailed in Appendix C. Data is from the ATUS, 2013-2017.

There are two issues that are important to mention when counting home production and childcare as non-leisure activities. The first is the assumption that home production has the same disutility as market work. Ramey and Francis (2009) provide evidence that people find some home production activities more pleasurable than market work. Additionally, there is an argument as to whether care for children and other family members should be included in home production since those services might be difficult to purchase on the market. The second matter is that home production in our framework only accrues a cost with no offsetting benefit. That is, the utility function does not include "home consumption." The reason for this is we would need to take a stand on the home production technology. While both these concerns are valid, we think our approach is useful because it provides an upper bound on the effects of home production. By making home production and childcare as displeasurable as market work, our measure of leisure, and women's leisure in particular, is most affected. We later show in Section 5 than counting home work as leisure widens the welfare differences between men and women. Counting home production and childcare as work, therefore, gives us a more conservative estimate of welfare differences between sexes.

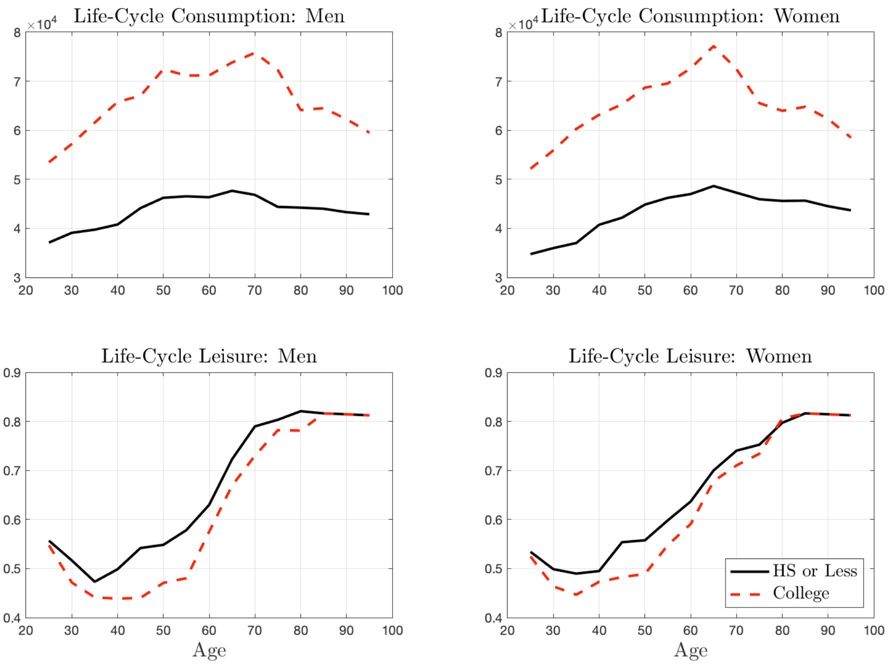

Figure 2: *Life-cycle Consumption Profiles.*

Notes: Top panels show average consumption by age and highest degree completed for men and women. Bottom panels show leisure as a share of time by age and highest degree for men and women. Consumption data is from the CE and leisure data is from the ATUS, 2013-2017.

Figure 2 summarizes life-cycle consumption and leisure profiles from 2013-2017. Consumption is always higher for those with more education. For most groups, consumption increases at the beginning of the life cycle and then drops off during retirement. This is broadly consistent with Fernndez-Villaverde and Krueger (2007) and Gourinchas and Parker (2002). The life-cycle trends in leisure are shown in the bottom two panels of Figure 2. These trends are also broadly consistent with the evidence. Leisure is higher for groups with less education as observed in Aguiar and Hurst (2007) and leisure increases towards retirement age, consistent with French (2005).

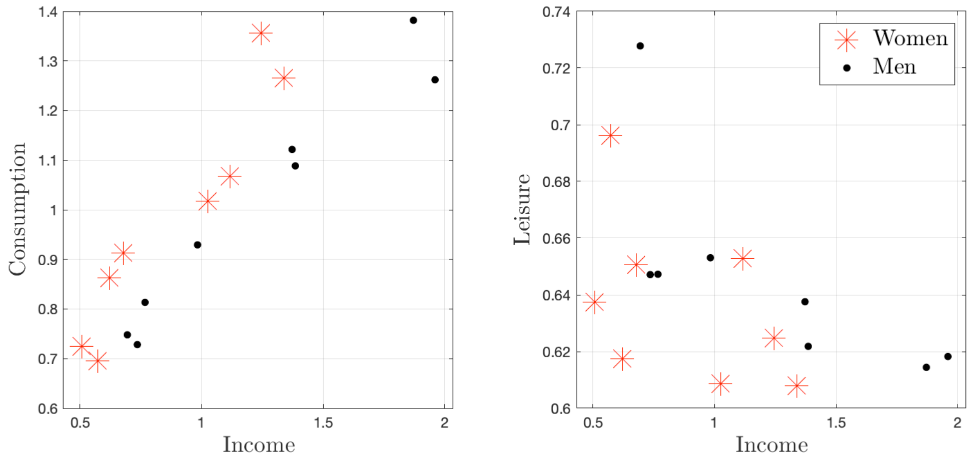

Figure 3: *The Relationship between Consumption and Leisure with Income.*

Notes: Left panel plots per person consumption by group against income and right panel plots per person leisure against income. Income and consumption data is from CE and leisure data is from the ATUS, 2013-2017.

To show how consumption and leisure are related to income, we obtain income data from the CE. We calculate per capita income by summing each individual's pre-tax wage and self-employment income with the household's other sources of income divided by the number of adults in the household. As an example, if a household consists of two members, A and B, who earn \$40,000 and \$50,000 in wage income respectively and collectively the household earns \$10,000 in dividend income, we would assign \$45,000 in income for A and \$55,000 for B. The panel on the left of Figure 3 shows how lifetime income varies with lifetime consumption across groups[9] where income and consumption are normalized by their respective US averages. The group observations are further color coated by sex. As one would expect, consumption and income are positively correlated with a correlation coefficient of about 0.89. Meanwhile, the right panel of Figure 3 shows that income and leisure are negatively correlated with a correlation coefficient around -0.53.

[9]By 'lifetime' we mean the present discounted value of each flow measure. We use a discount rate of two percent per year.

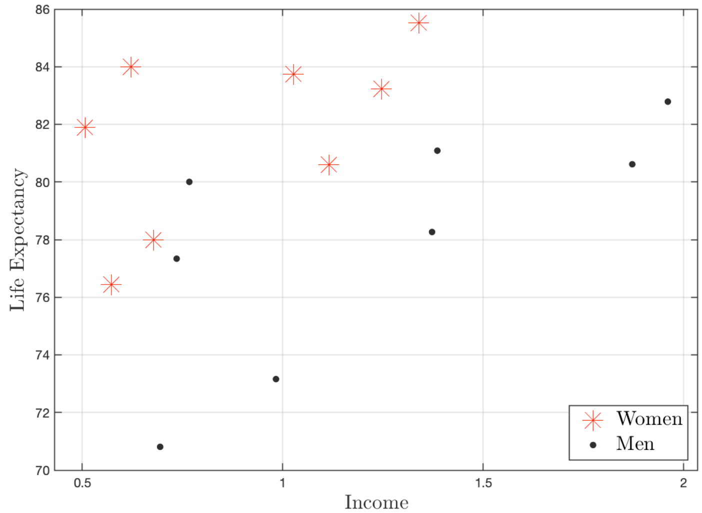

Figure 4: *The Relationship between Life Expectancy and Income.*

Notes: This plots per person (normalized) lifetime income by demographic group against life expectancy at age 25. Data is from 2013-2017.

The final data series is mortality rates. We get the number of deaths by age for each demographic group from the US National Vital Statistics System of the National Center for Health Statistics. We divide the number of deaths for each demographic group by the population of that group in the American Community Survey (ACS).[10] As in the CE and ATUS, we pool all the years 2013-2017 and average over five-year age bins. Once we have mortality rates we can convert them into life expectancies at age 25. The relationship with income and life expectancy is shown in Figure 4. Income and life expectancy are modestly correlated with a correlation coefficient of about 0.35. This correlation of income and life expectancy within genders is 0.69 and 0.56 for men and women respectively. There is wide variation in life expectancies with Asian women (of all education levels) expecting to live into their mid 80s whereas the White and Black men without a bachelor's degree only expect to live until their lower 70s. To put this in perspective, the difference in life expectancy at birth between the United States and Cambodia is about 15 years.[11]

To summarize, income is strongly positively correlated with consumption, negatively

[10]Both the ACS data and the time use data were obtained in the IPUMS (Ruggles et al., 2020).

[11]Data is from https://www.cia.gov/library/publications/the-world-factbook/rankorder/2102rank.html.

correlated with leisure, and modestly positively correlated with life expectancy. Because lifetime utility is increasing in consumption, leisure, and life expectancy, these empirical relationships point to ambiguous effects on overall welfare, which, again, motivates our exercise.

3.2 Model parameter values

While we pool our results into five-year age bins, the model period is one year. For example, we pool the data for white college-educated men from 25-29 to calculate consumption, leisure, and flow utility. So in the model, White college-educated men from 25-29 will have the same values for these variables. Following Jones and Klenow (2016), we set the discount rate, β, to 0.99. We set the growth rate in annual consumption to 1.25 percent which was the rate of growth in real expenditures on nondurables and services between 2013 and 2017 in the NIPA.

Recall, the within period utility function is $\bar{u} + \ln c + v(l)$. We assume $v(l) = -\frac{\theta}{1+\frac{1}{\epsilon}}(1-l)^{1+\frac{1}{\epsilon}}$ where $\theta > 0$ and $\epsilon > 0$. The intratemporal first order condition for leisure choice can be written as

$$\theta(1-l)^{\frac{1}{\epsilon}} = (1-\tau)\frac{w}{c}.$$

In models where leisure is defined as the difference between the time endowment and market work, the Frisch elasticity is equal to ϵ. Because our definition of leisure accounts for home production and child care, the Frisch elasticity is no longer constant. We choose ϵ to target an average Frisch elasticity equal to one.[12] We let $\tau = 0.21$ so as to match the average labor income tax from 2013-2015 according to the series assembled by McDaniel (2007). Assuming a Cobb-Douglas production function with an elasticity of output with respect to labor equal to two thirds, we calibrate θ to be consistent with a consumption to output ratio of two thirds and average time worked in the ATUS. This implies a value for θ of 17.08.

[12] In particular, the Frisch elasticity equals $\epsilon\frac{h_{work}+h_{home}+h_{child}}{h_{work}}$. In calibrating ϵ, we use the average values of market work, home production, and child care.

We choose $\bar{u}$ so that the value of a statistical life for the average person at age 40 is equal to \$7,000,000 in 2012 dollars. This is conceptually identical to Jones and Klenow (2016) and Falcettoni and Nygaard (2020). To put the value of statistical life in utility units, multiply by the marginal utility of consumption at age 40. That is

$$V_{40} = \frac{\$7,000,000}{c(40)}.$$

So the value of life at age 40 is

$$\frac{\$7,000,000}{c(40)} = \sum_{a=40}^{99} \beta^{a-40} S(a)[\bar{u} + \ln c(a) + g(a-40) + v(l(a))].$$

Solving for $\bar{u}$ gives

$$\bar{u} = \frac{\frac{\$7,000,000}{\bar{c}(40)} - \sum_{a=40}^{99} \beta^{a-40} S(a)[\ln c(a) + g(a-40) + v(l(a))]}{\sum_{a=40}^{99} \beta^{a-40} S(a)}$$

This implies a value of $\bar{u} = -5.81$.[13] Of course, changing either the value of a statistical life or the utility function will imply a different $\bar{u}$. We consider both changes in the robustness section. Table 1 shows the parameter values.

Table 1: Parameter Values

Parameter	Interpretation	Value	Target
$\bar{u}$	Constant in utility function	-5.81	VSL of \$7,000,000 (2012 dollars)
ϵ	Curvature on leisure in utility function	0.54	Frisch elasticity equal to 1
θ	Disutility of labor	17.08	Average hours of 0.28
g	Annual Growth rate of consumption	0.0125	Growth of real consumption in NIPA
β	Annual Discount rate	0.99	Jones and Klenow (2016)

[13] Falcettoni and Nygaard (2020) calibrated value of $\bar{u}$ is 7.2. While their utility function is slightly different than ours, the main difference is that they normalize per capita consumption to be 1. Aggregate real per capita consumption for 2013–2017 is about \$35,200. Adding $\ln(35,200)$ to our value of $\bar{u}$ gives about 4.66 which is much closer to their value.

4 Results

Figure 5 shows the cross-sectional distribution of economic welfare, measured as the equivalent variation of each demographic group relative to the average American. A positive (negative) number means that the group is better (worse) off than the average American. Quite intuitively, welfare is generally increasing in educational attainment within gender and race/ethnicity categories. Recall from Figure 1 that within each racial and gender category, the higher educated devote less time to leisure than the less educated. This means that the additional consumption and life expectancy for highly educated groups dominates the disutility from more work. Within education and race-ethnicity categories, women have higher welfare than men. While woman and men have fairly equal consumption and leisure streams, women tend to live longer. Finally, within gender and education categories, Asians tend to have the highest welfare, Blacks the lowest, and Whites and Hispanics somewhere in the middle.

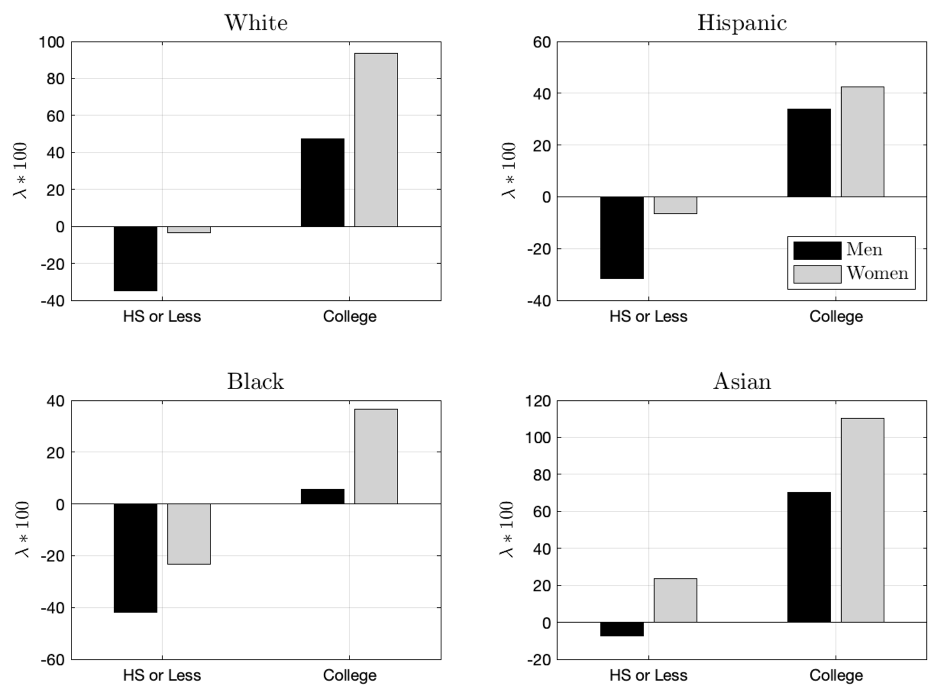

Figure 5: *Baseline Equivalent Variation.*

Notes: Equivalent variation is amount one would change the average American's consumption so that lifetime utility equals that of each demographic group. It is expressed as $\lambda \times 100$ where λ comes from Eq. (2). Time period is 2013-2017.

While the qualitative results provide a sense of the welfare orderings, we are ultimately after quantitative results. How big are the welfare differences quantitatively? Quite large. The economic welfare of the average college-educated Asian woman is 110 percent higher than the welfare of the average American while the average Black or White man without a bachelor's degree is around 40 percent lower. Translated into equivalent variations, a person would be indifferent between living the life path of the average Black or White man without a college degree and the life path of the average American with 40 percent lower consumption. Put another way, the average American would be willing to forgo up to 40 percent of annual consumption to avoid living the life path of a White or Black man who did not complete college.

On the other end of the distribution, the average American would need to raise their consumption by 110 percent to be indifferent between living that life path and the life path of the average college-educated Asian woman. These figures are also large for college-educated White women (94 percent) and college-educated Asian men (70 percent). The college education premium, defined in this context as the difference in welfare between those with and without a bachelor's degree, ranges from 48 percentage points (Black men) to 97 percentage points (White women) and is 82 percentage points on average. The numbers may represent an upper bound for the welfare premiums of additional school because we start the analysis after education decisions have been made and their resource and time costs have been incurred.[14]

Within each race/ethnic and education category, there is a clear gender gap whereby women have higher welfare than men. On average, welfare is 27 percent lower for men than women. To put this in comparison to more conventional gender gaps, such as the gender wage gap, the unconditional gender wage gap (i.e. the ratio between average annual earnings for women and annual earnings for men) in 2014 was about 79 percent (Blau and Kahn,

[14] We say "may represent" because there might be a consumption value to school that dominates the sum of the time cost and resource cost paid at the time. If college graduates in the data are paying back loans, that will be reflected in lower consumption.

2017). Consequently, our gender welfare gap is quantitatively similar to, but the opposite sign of, more conventional definitions of the gender gap.

How do cross-sectional differences in welfare compare to cross-sectional differences in income? Within each group we calculate average pre-tax wage income and total income weighting by the CE sample weights. Lifetime income is defined as the present discounted value of income over all ages for a given group with a discount rate of two percent. All of our measures of income are in relative terms, i.e. group i's income relative to the average American. In calculating the cross-sectional standard deviations, we weight each group by their share in the population according to the 2013-2017 ACS. Table 2 shows the results. The standard deviation of economic welfare is 11 percent higher than in lifetime income and is quantitatively similar to income inequality at midlife. The standard deviation of consumption is about half of the standard deviation of welfare. These results point to potential shortcomings of inferring welfare inequality from consumption inequality across demographic groups.

Table 2: Standard Deviations of Relative Income and Welfare

	St. Dev.
λ	12.11
Lifetime income	10.90
Average income	10.69
Average consumption	6.14
Age 40-54 income	12.84
Age 40-54 wage income	14.62

Figure 6 shows the correlation between lifetime income (expressed in relative terms) and economic welfare. There is an obvious positive, albeit imperfect, correlation between welfare and income. The correlation is about 0.64. This is much lower than the correlations between welfare and GDP across countries as in Jones and Klenow (2016) and across states as in Falcettoni and Nygaard (2020). Part of this is coming from the fact that, on average, women earn lower income over their lifetimes but enjoy longer life expectancies. The correlation between welfare and income within genders is 0.93 for women and 0.92 for men. Falcettoni and Nygaard (2020) report a correlation of about 0.80 across American states.[15]

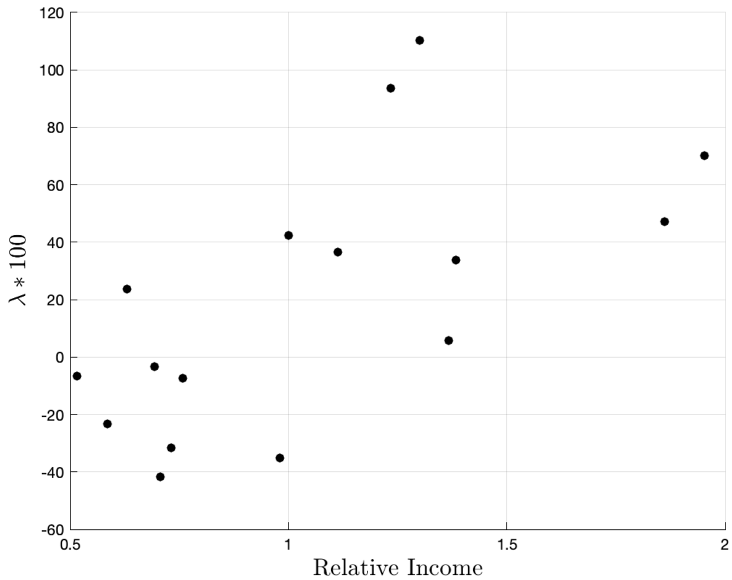

Figure 6: *Relationship between Lifetime Income (normalized by US average) and Welfare.*

Notes: Equivalent variation is amount one would change the average American's consumption so that lifetime utility equals that of each demographic group. It is expressed as $\lambda \times 100$ where λ comes from Eq. (2). Income by demographic group is normalized to 1. Time period is 2013-2017.

4.1 Decomposing the components of welfare

As we discussed in the previous section, economic welfare can be decomposed into its three constituent parts: life expectancy, average consumption, and average leisure. Figures 7 and 8 show the decomposition of the results for men and women, respectively. It is clear visually how big of a role consumption and life expectancy play. For instance, Black women with less than a bachelor's degree have 23 percent lower welfare than the average American. Leaving the average American's consumption and leisure the same, going from the life expectancy of the average American to the average less than college-educated Black woman (a difference of about 1.3 years) lowers welfare by about 10 log points or 9.78 percent

[15]In Appendix B we compare our results with Falcettoni and Nygaard (2020) by calculating welfare for each state based off demographic composition.

$(100 * [\exp(-0.1029) - 1])$. Likewise, the lower average consumptions for a less than college educated Black women lowers welfare by 36 log points and higher average leisure raises welfare by 20 log points.

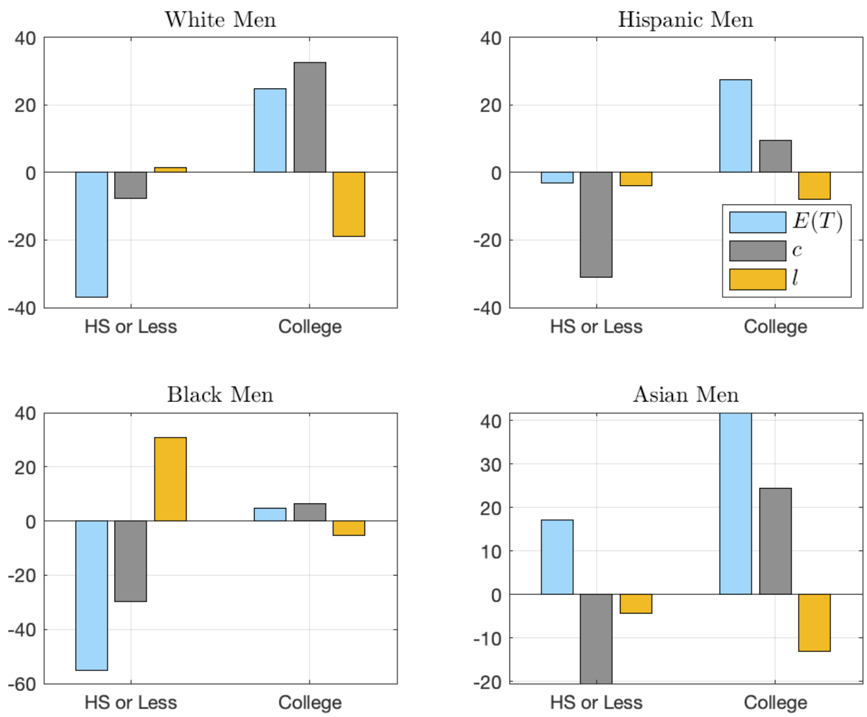

Figure 7: *Decomposition of Economic Welfare for Men.*

Notes: Equivalent variation is decomposed into its three components shown in Eq. (3). $E(T)$: life expectancy; c: consumption; l: leisure. Time period is 2013-2017.

With the exception of Asian and Hispanic women, the educational gradient of life expectancy is enormous. Hispanic men who graduate college expect to live nearly three years longer than those with less than a college degree. This raises welfare by about 35 percent. Similarly, White women who graduate college expect to live close to five years longer than those with less than a college degree. This raises economic welfare by 58 percent.

In summary, there is a high degree of welfare dispersion in the data. The standard deviation of welfare is comparable to the standard deviation of earnings in midlife and more than double the standard deviation of consumption. These vast differences are primarily driven by differences in life expectancy and consumption, with differences in leisure playing a less prominent role. Welfare is higher for women than men, highest for Asians and lowest for Blacks, and, unsurprisingly, increasing in educational attainment.

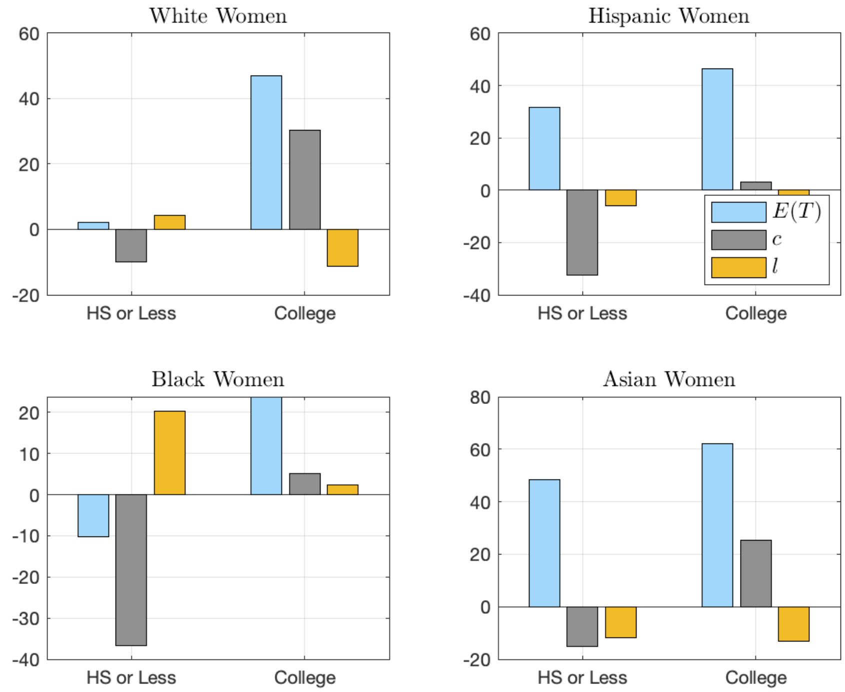

Figure 8: *Decomposition of Economic Welfare for Women.*

Notes: Equivalent variation is decomposed into its three components shown in Eq. (3). $E(T)$: life expectancy; c: consumption; l: leisure. Time period is 2013-2017.

4.2 The Evolution of Welfare over Time

We now turn our attention to how welfare has changed over time for each group. In particular, we calculate the percent one would have to change consumption of the average member of group i in 2000-2004 to be indifferent between living as a member of that group in 2013-2017 and living as a member in that group in 2000-2004 with the revised consumption profile. We choose the starting date of 2000 because that is the first year of the American Community Survey.[16] In principle, we could also calculate mortality rates using the 1980 or 1990 censuses, but the race and ethnicity classifications have changed over time which

[16] The ATUS starts in 2003, so in measuring welfare in early period, we use the time use data in 2003 and 2004.

complicates comparisons with the more current data.

Figure 9 shows the results. In the figure, a group experienced welfare gains if the equivalent variation is positive. With the exception of Asian men with less than a college degree, each demographic group experienced welfare gains over the 14-year interval. Within each racial/ethnic group, the college educated experienced the highest welfare gains. Black men of all education categories experienced welfare gains of about 50 percent, or about three percent per year. Welfare gains are also higher than three percent annually for Black and Hispanic women with a college degree. The lowest gains over this period were for Whites and Asians with less than a college degree. The correlation between welfare growth and initial welfare levels is -0.26, implying a limited degree of convergence.

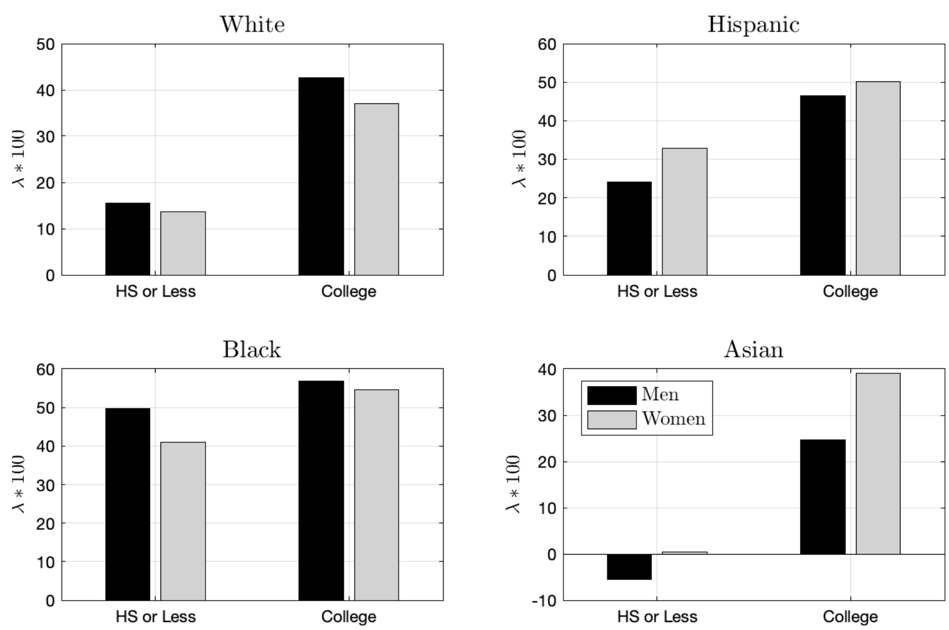

Figure 9: *Welfare Growth 2000-2004 to 2013-2017*

Notes: Welfare growth is the equivalent variation is amount one would change the average member of each demographic group's consumption in 2000-2004 so that lifetime utility equals that of the same demographic group in 2013-2017.

To understand the components driving economic welfare growth within each demographic group, we decompose the welfare gains into the contributions of changes in life expectancy, consumption, and leisure. The results are plotted in Figures 10 and 11 for men and women, respectively. The increase in life expectancy positively contributed to welfare growth across all demographic groups and groups with the highest welfare growth also experienced some of the largest gains in life expectancy. Black men of all education levels, for instance, had the largest reductions in mortality. The life expectancy at age 25 for Black men without a college degree increased from 67.1 to 70.8 years and from 73.8 to 78.3 for those with a college degree. The lowest gains in life expectancy are for White men and women without a college degree.

While consumption increased for all groups but one, the contribution of decreased mortality is more quantitatively significant for most groups than changes in consumption or leisure. Thus, this decomposition highlights the distinction between quantity of life (more life) versus quality of life (more instantaneous utility) for our welfare growth calculation. In this time period, the quantity of life is a key driver of welfare growth for many of the demographic groups.

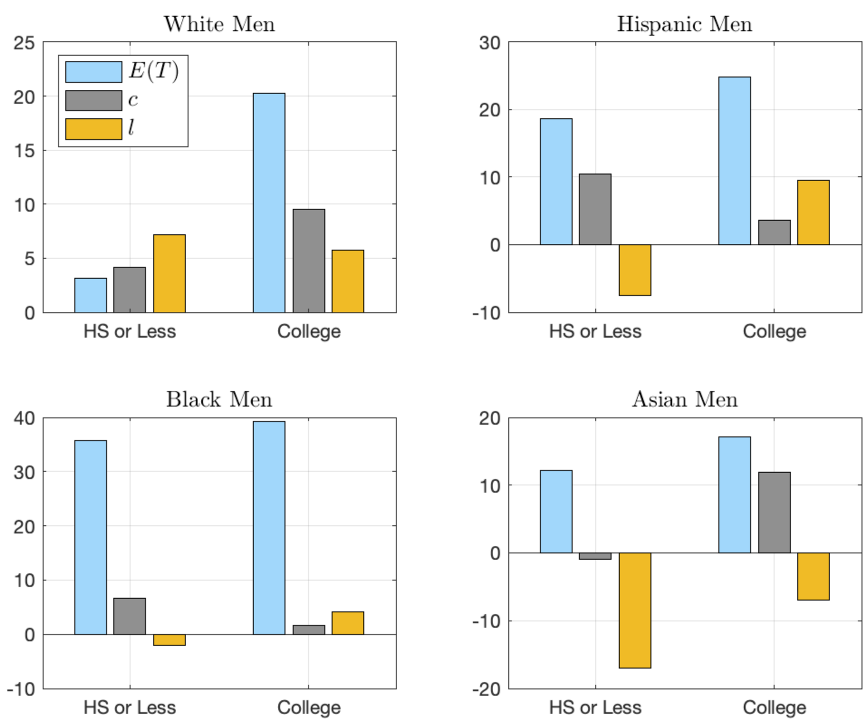

Figure 10: *Decomposition of Economic Welfare for Men 2000-2004 to 2013-2017*

Notes: Equivalent variation is decomposed into $E(T)$: life expectancy; c: consumption; l: leisure.

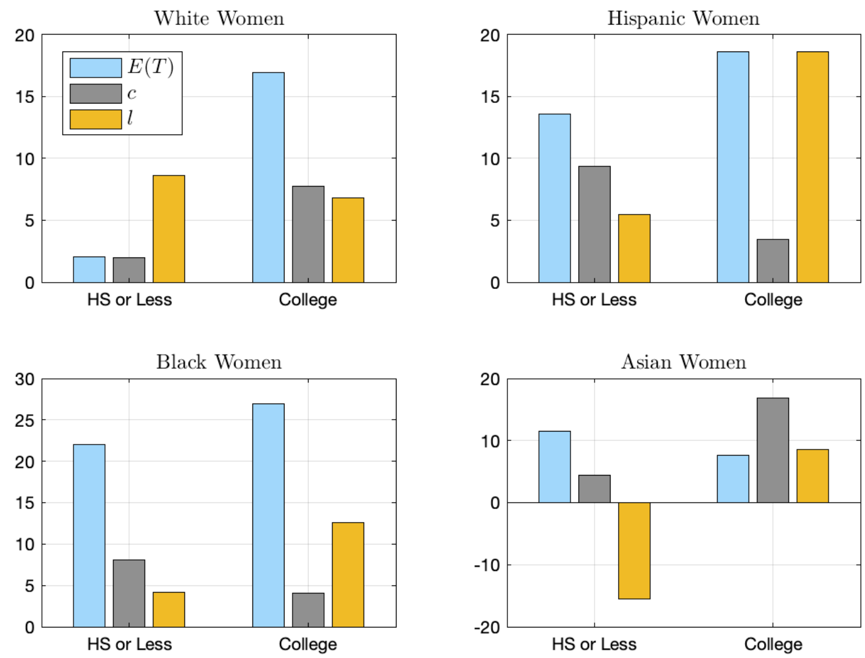

Figure 11: *Decomposition of Economic Welfare for Women 2000-2004 to 2013-2017*

Notes: Equivalent variation is decomposed into $E(T)$: life expectancy; c: consumption; l: leisure.

Although any implications for policy are limited by this being a pure measurement exercise, our results show the large potential gains to reducing mortality rates for some groups. Moreover, any policy that affects mortality rates will have distributional implications. Our framework can be used in equilibrium models to compare the costs, benefits, and distributional consequences of policies that improve the quality of life (e.g. by raising consumption) or extend the quantity of life.

Welfare Growth and Education Composition: The welfare growth calculations isolate changes within each gender/race/education group. Over time, however, educational attainment has increased across both genders and all race/ethnic groups. What we want to know next is how the expected utility of each gender and racial/ethnic categories is affected by the greater share of each population graduating college. For each race/ethnicity and gender we represent expected utility as

$$\mathbb{E}\left[U_t\right] = \pi_{c,t}U_{c,t} + (1 - \pi_{c,t})U_{hs,t}$$

where π_c is the share of college graduates for the gender-race/ethnicity group, U_{hs} is the lifetime utility of someone in the group with less than a college degree and U_c is the lifetime utility of someone in the group with at least a college degree. Thus, a group's expected utility can increase between 2000-2004 and 2013-2017 because either the lifetime utility for one or both education groups increase or because the share of college graduates increase. We calculate the following expected utilities

$$\begin{aligned}
\mathbb{E}\left[U_{2004}\right] &= \pi_{c,2004}U_{c,2004} + (1 - \pi_{c,2004})U_{hs,2004} \\
\mathbb{E}\left[U_{2017}\right] &= \pi_{c,2017}U_{c,2017} + (1 - \pi_{c,2017})U_{hs,2017} \\
\mathbb{E}\left[\widehat{U}_{2017}\right] &= \pi_{c,2004}U_{c,2017} + (1 - \pi_{c,2004})U_{hs,2017}.
\end{aligned}$$

That is, we compare expected utilities in the early period to expected utilities in the late period and a counterfactual expected utility where we hold the college shares fixed at their 2004 level. Welfare growth is defined as the change in 2004 consumption required to equate expected utility in 2004 to expected utility in 2017.

Table 3 summarizes these results for each gender-race/ethnicity combination. The first two columns show the share of the population by gender with a college degree in the 2000-2004 period and the 2013-2017 period. The column titled "Base" is the equivalent variation between expected utility in 2004 ($\mathbb{E}\left[U_{2004}\right]$) and 2017 ($\mathbb{E}\left[U_{2017}\right]$). The column "Counter" is the equivalent variation between expected utility in 2004 and the counterfactual expected utility in 2017 ($\mathbb{E}\left[\hat{U}_{2017}\right]$) which holds the college shares fixed at their 2004 level.

Table 3: Welfare Change and Education Composition

	Men				Women			
	College Share		Eq. Var		College Share		Eq. Var	
	00-04	13-17	Base	Counter	00-04	13-17	Base	Counter
White	30.09	40.48	35.19	23.63	25.87	39.45	31.08	19.44
Black	15.45	24.28	60.61	50.97	16.60	29.48	54.95	43.25
Asian	49.60	64.74	18.83	8.63	42.09	61.95	27.23	15.40
Hispanic	11.39	23.33	37.58	26.61	11.86	26.81	43.33	34.89

The difference between the "Base" and "Counter" equivalent variations can be attributed to improvements in education composition between the early and late period. In percent terms this difference is the biggest for Asian men. This is because Asian men without a college degree experienced a decline in welfare so all the growth in expected utility comes from increased welfare of the college educated and improvements in composition.

A counter-intuitive result is that while the welfare gains of Black men for both education group from 2000-2004 to 2013-2017 are around 50 percent (see Figure 9), the gain to expected utility is around 60 percent. This is a case when the total is greater than the sum of its parts. Because welfare is higher for college-educated Black men than Black men without college, an increase in the share of college educated increases expected utility even if there was no change in welfare within education groups. Adding to this the fact that welfare *did change* within education groups (rising by about 50 percent) implies that the percent change in average welfare for Black men was bigger than the percent change for Black men with a college degree. In this way, the welfare growth by individual education levels understates the total welfare growth for Black men due to compositional changes in education. This is the only group where this occurs.

These results show that even though some of the changes over time are driven by composition, the fact remains that there is increasing dispersion along the dimension of education. The disparity of welfare gains over the last two decades could be connected to the political polarization that Autor et al. (2020), among others, speak to. On the other hand, the small degree of convergence between groups that we observe is similar in spirit to the converging occupation shares across demographic groups that Hsieh et al. (2019) highlight. Their analysis starts in 1960 when there was much more misallocation, so it's natural to expect less convergence in our sample years.

Welfare Growth and Deaths of Despair: Next, we focus on a change over the last 20 years that has acutely affected mortality rates, and therefore economic welfare, for some demographic groups. Case and Deaton (2017) document a rise in midlife mortality for less educated White Non-Hispanics beginning around 1998. They trace these increasing mortality rates to the proliferation "deaths of despair" which include suicide, drug and alcohol poisoning, chronic liver disease, and cirrhosis. Indeed, the increase in deaths of despair were large enough to increase mortality rates despite mortality rates of other causes of death declining. They conclude that had mortality rates stayed at their 1998 level, 96,000 deaths would have been avoided between 1999–2013. In this section we ask what would happen to economic welfare had the probability of a death of despair in 2013-2017 stayed the same as 2000-2004.

Our quantitative exercise is slightly different from the earlier sections. In particular, we ask how much consumption a person in a given demographic group facing 2013-2017 mortality rates (and 2013-2017 levels of consumption and leisure) would be willing to sacrifice to live an alternative life where the probability of a death of despair equals its value from 2000-2004. Our data includes the cause of death, which allows us to measure the contribution of deaths of despair to overall mortality. In this exercise, we compute counterfactual mortality rates for 2013-2017 by replacing the mortality rates from deaths of despair with their 2000-2004 levels and leave the death rate from all other causes equal to their 2013-2017 levels. While we proceed with the data on cause of death as reported, there may be reporting errors in the exact cause of death. For example, McGivern, Shulman, Carney, Shapiro, and Bundock (2017) find that 51 percent of a sample of 601 death certificates from the Vermont Electronic Death Registration System had major errors in their classification. Although it is unclear how this translates to classification error rates across the country, we do acknowledge that our results may be sensitive to misclassifications in cause of death.

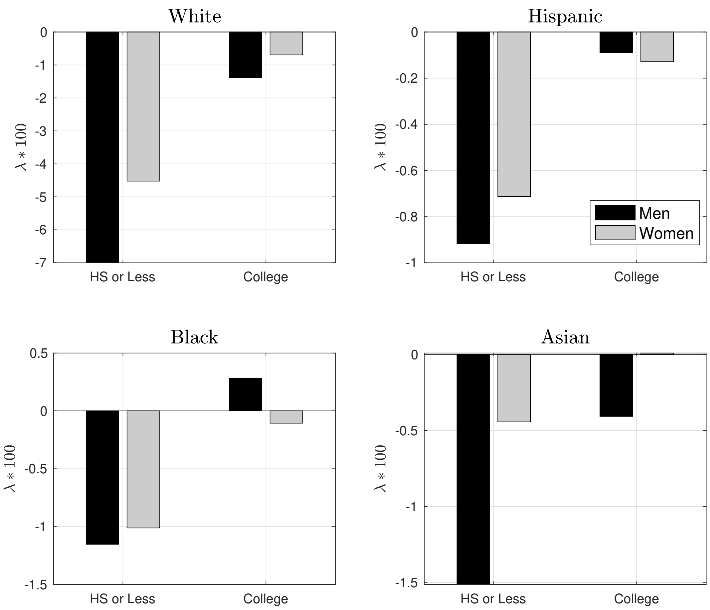

Figure 12: *Welfare Gains from Counterfactual Deaths of Despair*

Notes: Deaths of despair are defined as in Case and Deaton (2017). Counterfactual deaths of despair replaces mortality rates from deaths of despair for each demographic group in 2013-2017 with each group's deaths of despair observed in 2000-2004. Equivalent variation is amount one would change the average demographic group's consumption in 2013-2017 so that lifetime utility equals that of each demographic group facing counterfactual deaths of despair, holding all consumption and leisure components constant.

Figure 12 shows the results. Consistent with the findings of Case and Deaton (2017), the welfare losses are largest for White men and women. A White man without a college degree would give up about seven percent of annual consumption to return to the 2000-2004 mortality rates for deaths of despair. At the other end of the distribution, college-educated Black men experienced a decrease in deaths of despair between 2000-2004 and 2013-2017.

By and large, people with less education would be willing to sacrifice more consumption to avoid deaths of despair. This is particularly true among Whites. Our findings are consistent with Novosad, Rafkin, and Asher (2020) who show that increasing mortality rates among Whites is attributed to rising mortality in the bottom part of the education distribution. Referring back to the decomposition exercise of the sources of welfare growth (Figures 10 and

11), we showed that less educated White men and women had the lowest welfare gains from increases in life expectancy from 2000-2004 to 2013-2017, and the rise of deaths of despair contributes to this finding.

5 Robustness

We consider a number of robustness exercises and alternative specifications to understand the importance of various factors in our main results. First, we reclassify home production and childcare as leisure activities. This allows us to consider alternative definitions of work and leisure. Second, we decompose the importance of within-group inequality for our results. We then consider heterogeneity in parameter values based on educational attainment. We also generalize the utility function and calculate results across a range of parameter values and consider alternative definitions of consumption, different price deflators, and different parameter values in the instantaneous utility function. Finally, we calculate compensating variations as an alternative to equivalent variations.

5.1 Home Production as Leisure

The baseline specification categorized home production and childcare as non-leisure time, treating these components the same as market work. In this section we count home production and childcare as leisure. As we showed in Figure 1, men, on average, spend more time on market work relative to women while women spend more time on home production and childcare than men. The result of this exercise, therefore, will be to endow women with more leisure time than men.

For this exercise, we use the CE in place of the ATUS for time spent on market activities. All of the other aspects of the model and data inputs remain the same as the baseline specification. The reason we use the CE in place of the ATUS is a practical one: we need these results for the CE to use as a basis for comparison to our next robustness exercise and the results using the CE and the ATUS without home production are virtually the same.

The gray series in Figures 13 and 14 shows the results for men and women, respectively. The black series is the baseline results for comparison (we discuss the last series in the next section). The economic welfare of men decreases relative to the baseline when home production is counted as leisure. This is intuitive since we are decreasing the leisure of men relative to the leisure of the average American. On the other hand, the results for women – with the exception of Black women – in Figure 14 show an increase in welfare. Recall that the economic welfare of men was 27 percent lower than women in our baseline result. Counting home production and childcare as leisure increases the gender gap to 43 percent. The results from counting home production and childcare as leisure, therefore, can be interpreted as an upper bound of the average welfare differences between men and women.

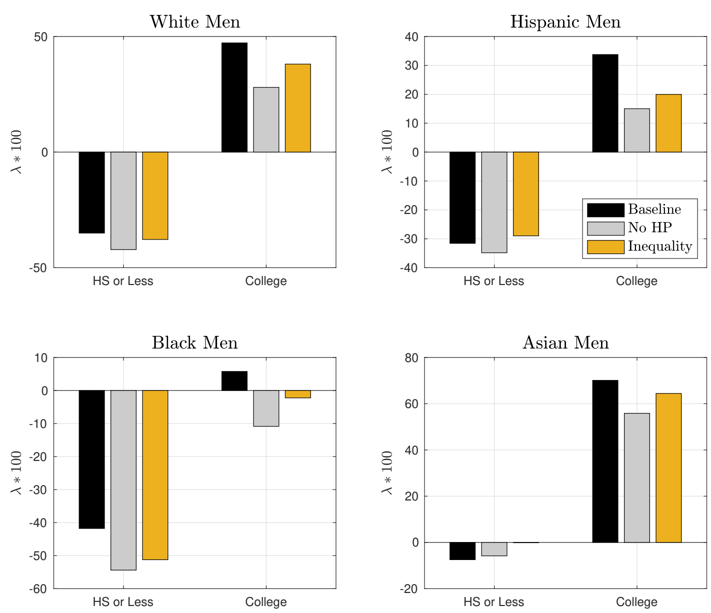

Figure 13: Economic Welfare for Men: Baseline, Without Home Production, and Inequality

Notes: The black series is the baseline results. The middle series is the model classifying home production as leisure. The right series is welfare calculated by Equation 3.

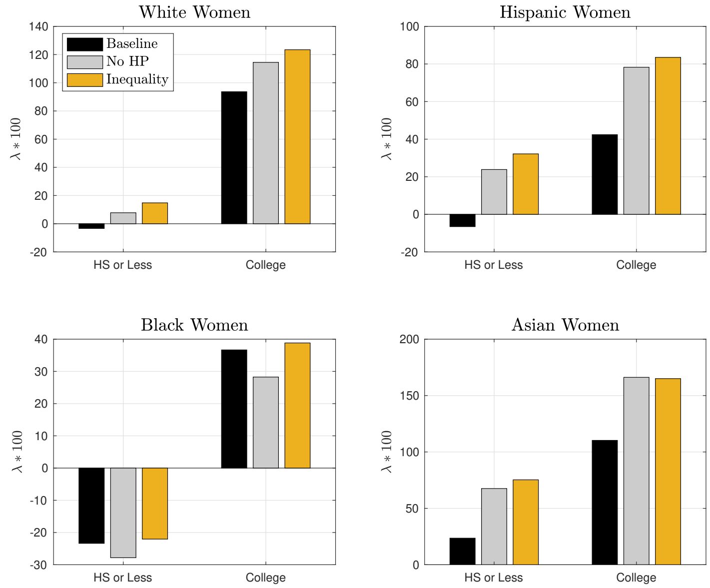

Figure 14: Economic Welfare for Women: Baseline, Without Home Production, and Inequality

Notes: The black series is the baseline results. The middle series is the model classifying home production as leisure. The right series is welfare calculated by Equation 7.

5.2 Within-Group Inequality

In this section, we isolate the role of within-group inequality on welfare. To do this, we use the accounting exercise from Jones and Klenow (2016). In this exercise, we again use the CE consumption and leisure data (without home production) which allows us to get the joint distribution of consumption and leisure. Here, the conceptual idea changes slightly. Just as in the baseline, we are looking at lifetime utility from the perspective of someone at age 25 who is assigned a particular demographic group. But now, each individual factors in consumption and leisure possibilities they may receive as being a member of this group. For example, if there was no inequality in consumption and leisure within a demographic group, welfare would be evaluated at the means for both variables, the same as the baseline. With inequality in consumption and leisure, concavity in the utility function results in lower mean utility than utility evaluated at the means. That is, the prospect of drawing low consumption or leisure dampens the level of utility for a particular group.

As before, we index each gender, race, and education level combination by i and age by a so there are $N_i(a)$ people in an age demographic of a particular group. Letting j denote a member of $N_i(a)$, lifetime expected utility is

$$\bar{V}_i = \sum_{a=25}^{99} \beta^{a-25} S_i(a) \sum_{j=1}^{N_i(a)} \omega_{i,j}(a) \left[\bar{u} + \ln c_{i,j}(a) + g(a-25) + v(l_{i,j}(a))\right] \tag{6}$$

where $\omega_{i,j}(a)$ is the sample weight of person i in group $N_i(a)$. All of the remaining notation is identical to the baseline.[17] With additively separable utility and common weights for consumption and leisure ($\omega_{i,j}(a)$), we can decompose economic welfare in this case into five different components:

$$\begin{aligned}\ln(1+\lambda_i) = &\underbrace{\sum_a \Delta s_i(a) u_i(a)}_{\text{Life expectancy}} + \underbrace{s_{US}(a)(u(\bar{c}_i) - u(\bar{c}_{us}))}_{\text{Consumption}} + \underbrace{s_{US}(a)(v(\bar{l}_i) - \bar{v}(\bar{l}_{us}))}_{\text{Leisure}} \\ &+ \underbrace{s_{US}(a)([\bar{u}(c_i) - u(\bar{c}_i)] - [\bar{u}(c_{us}) - u(\bar{c}_{us})])}_{\text{Consumption inequality}} + \underbrace{s_{US}(a)(\left[\bar{v}(l_i) - v(\bar{l}_i)\right] - \left[\bar{v}(l_{us}) - v(\bar{l}_{us})\right])}_{\text{Leisure inequality}}\end{aligned} \tag{7}$$

where $\Delta s_i(a)$ and $s_{US}(a)$ are defined as in Equations 4 and 5. The derivation of Equation 7 follows that in Jones and Klenow (2016) (see Equation (19) in their paper).

The results of this are the last series in Figures 13 and 14 which shows the results for men and women, respectively. The basis of comparison is the gray series, "No HP." This is the same data evaluated at the means and the only differences between these series are the consumption and leisure inequality terms. The welfare measure is higher in all cases when inequality is measured. The intuition for this comes from the fact that "within-group" inequality for the average American is simply total inequality which is going to be bigger than "within-group" inequality defined by more narrow race, gender, and education combinations.

[17] We normalize the sample weights so that $\sum_{j=1}^{N_i(a)} \omega_{i,j}(a) = 1$ for all (i,a) pairs.

Moreover, the results are not substantially different when inequality is not included. Thus, inequality between groups is more important quantitatively than inequality within groups.

The fact that within-group inequality does not change our results much is reassuring because we are unable to compute within-group inequality in our baseline. The reason for this is because the CE and ATUS have different sample weights. The decomposition in Equation 7 relies on having identical weights, $\omega_{i,j}(a)$. Since the CE does not include time use in home production and childcare, there is a tradeoff between capturing a more realistic definition of leisure and isolating within-group inequality. Given that within-group inequality plays a relatively small role quantitatively, we chose to use the more expansive definition of work in our baseline measure.

5.3 The College Premium

Starting our analysis at age 25 allows us to treat education as exogenous since the choice has already occurred. We also assume that the utility functions are identical irrespective of educational attainment. However, it is possible that preference heterogeneity may play a role in determining who attains a college degree and these difference in preferences could, in turn, affect welfare later in life. In this section we explore the possibilities of having different parameter values based on educational attainment. The welfare concept is slightly different than in the baseline case. Specifically, the equivalent variation in now calculated as the college welfare premium – the percent increase in the consumption of the less educated to equate their lifetime utility to the lifetime utility of a college graduate but within same gender/race group.

The results are shown in Table 4. For the baseline case, the college welfare premium for men is 127 percent and 90 percent for women. The rows below the baseline case consider the alternative parameter values. We note that these parameter choices are not calibrated, but rather chosen to highlight how different parameter values between education groups affect the main results. First, we consider lowering the subjective discount factor β from 0.99 to 0.985

for only those with less than a college degree, which implies that college-educated individuals are relatively more patient. The welfare results are most sensitive to this change. Even a relatively small change in β implies large changes to the cumulative effective discount rate, which has large effects on economic welfare. Next, motivated by the fact that college-educated individuals take less leisure than individuals without a college degree, we impose a 20 percent reduction in θ for the college educated. It follows that the college premium increases in this scenario. Likewise, the college premium increases if the consumption growth rate g_c is higher for the college educated. We do note that in the data (see Figure 2) the consumption profiles for the college educated are steeper than the less college educated. So, in a way, the model already factors in higher consumption growth rates for the college educated irrespective of an increase in g. Finally, we allow $\bar{u}$, the constant in the utility function, to vary by education. Flow utility increases with $\bar{u}$ so welfare is relatively higher for the group with the higher $\bar{u}$.

Table 4: Welfare College Premium

	Men	Women
Baseline	127.52	89.75
β_{hs}=0.985	363.57	299.99
θ_c=0.8×θ	168.81	118.95
g_c =0.02	182.35	135.11
$\bar{u}_c$=-5.81+0.25	200.65	147.95
$\bar{u}_c$=-5.81-0.25	72.18	45.22

5.4 Other Robustness

In this section we show that our quantitative findings are robust to changing preferences, different ways of defining and measuring consumption, and changing the way we calculate welfare. All of the results in this section is in Table 5. For each exercise, we report the standard deviation of welfare across groups and the correlation of each exercise with the baseline results.

Table 5: Alternative Specifications

	St. Dev.	Correlation
Baseline	12.12	
σ=2	10.48	96.19
σ=0.5	12.31	99.09
Frisch=0.5	11.18	99.70
Frisch=2	12.45	99.97
VSL=$6 mil.	10.84	99.90
VSL=$8 mil.	13.45	99.93
g=0.5	11.74	99.99
g=2.0	12.5	99.99
Durables	12.67	99.57
CPI	12.08	99.99
PCE	12.15	99.99
Comp. Var.	11.4	99.81

First, consider a more general within-period utility function

$$u(c,l) = \frac{\left(c \exp\left(\frac{-\theta}{1+\frac{1}{\epsilon}}(1-l)^{1+\frac{1}{\epsilon}}\right)\right)^{1-\sigma} - 1}{1-\sigma}.$$

Our utility function in the previous sections was just the special case of $\sigma = 1$. We change the value of σ to 2 and 0.5 and recalibrate $\bar{u}$ and θ to match the VSL and average hours worked respectively. A higher value of σ lowers the dispersion of welfare and also leads to the lowest correlation with the baseline results. This is because more curvature in the utility function dampens the differences in instantaneous utility across groups, particularly those with higher welfare to begin with. Likewise, a lower value of σ exacerbates the differences in instantaneous utility, and hence the standard deviation of welfare.

The model results are not very sensitive to varying the Frisch Elasticity, which we get by changing ϵ, and welfare has a high degree of correlation with the baseline. Although the correlations with the baseline is high, higher Values of a Statistical Life (VSL) – $8 million compared to $7 million in the baseline – leads to more dispersion in the equivalent variations. This is not surprising in light of our findings that life expectancy accounts for much of the welfare differences across groups. As the VSL increases, those with low mortality rates benefit disproportionately. Since those with low mortality rates had high welfare to begin with, larger VSLs increase cross-group inequality. Higher consumption growth rates g also amplify welfare differences across groups. Again, those with higher consumption levels also have the longest life expectancies, on average. A higher consumption growth rate benefits those with longer life expectancies and higher consumption the most.

We also consider several variants of the consumption measure. In all of our previous results, we have used consumption of non-durables and services where each consumption category is deflated using a good-specific price deflator from the CPI. The first alternative consumption measure is to include durables, while continuing to deflate by the good-specific CPI. The welfare differences increase slightly. This is because the largest consumers of non-durable goods also tend to consume more durable goods as well, amplifying the consumption differences across groups. Next, we aggregate total non-durable consumption and deflate it all by the CPI and the PCE. This does not change the relative differences in consumption much at all and the results are very close to the baseline.

Finally, we compare the equivalent variations to compensating variations, i.e. the fraction of consumption the average member of group i needs to be compensated in order to be indifferent between living as a member of group i and the average person in the US. These two numbers will be different because groups do not share the same mortality rates. Nevertheless, the compensating and equivalent variations are very close to each other.

6 Conclusion

A wide range of economic research is devoted to understanding economic inequality in the US. We contribute to this body of research by documenting how economic welfare is distributed across the population and how that distribution has changed over the last 20 years. We find that, conditional on other demographic information, women have higher welfare than men, Asians have the highest welfare and Blacks the lowest, and welfare is rising with educational attainment. The qualitative rankings are robust to a variety of ways in defining consumption and leisure and parameter choices in the utility function.

We also make some progress towards understanding how income correlates with economic well-being. The correlation between our baseline measure of economic welfare and pre-tax lifetime income is about 0.6, but rises significantly when we compute the correlation within genders. In terms of dispersion, the standard deviation of welfare is 11 percent higher than the standard deviation of lifetime income and about double the standard deviation of lifetime consumption. While quality of life (which is a function of consumption and leisure) is distributed more equally than income, the vast differences in quantity of life, and the fact that quality and quantity of life are positively correlated, imply that the standard deviation of welfare is similar to that of income.

While we included several important components of welfare, we omitted innumerable others. A non-exhaustive list includes: utility from public goods, exposure to pollutants, morbidity, and occupational differences in the disutility of work. In terms of differential morbidity, Hosseini et al. (2019) develop a frailty index to study how health evolves over the life cycle. Kaplan and Schulhofer-Wohl (2018) document how feelings at work vary by race, education, and gender. We conjecture that our framework could be adapted to take into account both of these features. Thus, we see them as attractive areas for future research. Finally, Brouillette et al. (2021) calculate the disparate racial impact of Covid-19. One could calculate disparities across gender and education groups using our framework.

Acknowledgements

We are grateful to John Steenrod for research assistance on an earlier draft of this project. We also thank Greg Casey, Kevin Donovan, and Steven Lugauer for useful comments and seminar participants at the Central Bank of Uruguay, CEMFI, University of California Santa

Barbara, Amherst College, Osaka University, University of Massachusetts Lowell, the Virginia Association of Economists, University of New Hampshire, the Liberal Arts Macroeconomic Workshop, and the conference participants at the Welfare and Inequality in the 21st Century. The usual disclaimer applies.

References

Acemoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in Handbook of Labor Economics, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4 of Handbook of Labor Economics, chap. 12, 1043–1171.

Aguiar, M. and M. Bils (2015): "Has Consumption Inequality Mirrored Income Inequality?" American Economic Review, 105, 2725–56.

Aguiar, M. and E. Hurst (2007): "Life-Cycle Prices and Production," American Economic Review, 97, 1533–1559.

——— (2009): The Increase in Leisure Inequality, 1965-2005, Washington D.C., American Enterprise Institute for Public Policy Research.

Autor, D., D. Dorn, G. Hanson, and K. Majlesi (2020): "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure," American Economic Review, 110, 3139–83.

Bee, A., B. D. Meyer, and J. X. Sullivan (2015): "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?" University of Chicago Press, vol. 74, 204.

Blau, F. D. and L. M. Kahn (2017): "The Gender Wage Gap: Extent, Trends, and Explanations," Journal of Economic Literature, 55, 789–865.

Brouillette, J.-F., C. I. Jones, and P. J. Klenow (2021): "Race and Economic Well-Being in the United States," National Bureau of Economic Research Working Paper.

Case, A. and A. Deaton (2017): "Mortality and Morbidity in the 21st Century," Brookings Papers on Economic Activity, 48, 397–476.

Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016): "The Association Between Income and Life Expectancy in the United States, 2001-2014," Journal of the American Medical Association, 315, 1750–1766.

Falcettoni, E. and V. Nygaard (2020): "A Comparison of Living Standards Across the States of America," Finance and Economics Discussion Series 2020-041, Board of Governors of the Federal Reserve System.

Fernndez-Villaverde, J. and D. Krueger (2007): "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data," The Review of Economics and Statistics, 89, 552–565.

French, E. (2005): "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," Review of Economic Studies, 72, 395–427.

Goldin, C. and L. F. Katz (2008): "Transitions: Career and Family Life Cycles of the Educational Elite," American Economic Review, 98, 363–369.

Gourinchas, P. O. and J. A. Parker (2002): "Consumption Over the Life Cycle," Econometrica, 70, 47–89.

Heathcote, J., F. Perri, and G. L. Violante (2010a): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006," Review of Economic dynamics, 13, 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2010b): "The Macroeconomic Implications of Rising Wage Inequality in the United States," Journal of Political Economy, 118, 681–722.

Hosseini, R., K. A. Kopecky, and K. Zhao (2019): "The evolution of health over the life cycle," FRB Atlanta Working Paper 2019-12a.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019): "The allocation of talent and us economic growth," Econometrica, 87, 1439–1474.

Jones, C. I. and P. J. Klenow (2016): "Beyond GDP? Welfare across Countries and Time," American Economic Review, 106, 2426–2457.

Kaplan, G. and S. Schulhofer-Wohl (2018): "The Changing (dis-) Utility of Work," Journal of Economic Perspectives, 32, 239–58.

Krueger, D. and F. Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," Review of Economic Studies, 73, 163–193.

Liebler, C. A. and A. Halpern-Manners (2008): "A Practical Approach to Using Multiple-Race Response Data: A Bridging Method for Public Use Microdata," Demography, 45, 143–155.

Markides, K. S. and K. Eschbach (2005): "Aging, migration, and mortality: current status of research on the Hispanic paradox," The Journals of Gerontology Series B: Psychological Sciences and Social Sciences Special Issue 2, 60, 68–75.

McDaniel, C. (2007): "Average tax rates on consumption, investment, labor and capital," in Manuscript: Arizona State University.

McGivern, L., L. Shulman, J. K. Carney, S. Shapiro, and E. Bundock (2017): "Death Certification Errors and the Effect on Mortality Statistics," Public Health Reports, 132, 669–675.

Moss, M. (1973): The Measurement of Economic and Social Performance, vol. 38 of NBER Books, National Bureau of Economic Research, Inc.

Nordhaus, W. D. and J. Tobin (1971): "Is Growth Obsolete?" Cowles Foundation Discussion Papers 319, Cowles Foundation for Research in Economics, Yale University.

Novosad, P., C. Rafkin, and S. Asher (2020): "Mortality Change Among Less Educated Americans," Working Paper: Dartmouth University.

Ramey, V. A. and N. Francis (2009): "A Century of Work and Leisure," American Economic Journal: Macroeconomics, 1, 189–224.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2020): "IPUMS USA: Version 10.0 [dataset]," Minneapolis, MN: IPUMS, https://doi.org/10.18128/D010.V10.0.

Sandra L. Hofferth, Sarah M. Flood, M. S. and D. Backman (2020): "American Time Use Survey Data Extract Builder: Version 2.8 [dataset]," College Park, MD: University of Maryland and Minneapolis, MN: IPUMS, https://doi.org/10.18128/D060.V2.8.

Sasson, I. (2016): "Trends in Life Expectancy and Lifespan Variation by Educational Attainment: United States, 1990–2010," Demography, 53, 269–293.

A Derivations

A.1 Decomposition of the $\sigma = 1$ case

The present discounted value of lifetime utility when utility is log over consumption is given by

$$\bar{V}_i = \sum_{a=25}^{99} \beta^{a-25} S_i(a) \left[\bar{u} + \ln c_i(a) + g(a-25) + v(l_i(a)) \right].$$

Let $u_i(a) = [\bar{u} + \ln c_i(a) + g(a-25) + v(l_i(a))]$. The equivalent variation is defined implicitly by

$$\ln(1+\lambda_i) = \frac{\sum_{a=25}^{99} \beta^{a-25} \left(S_i(a)u_i(a) - S_{US}(a)u_i(a) + S_{US}(a)u_i(a) - S_{US}(a)u_{US}(a) \right)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)}$$

We can write the equation as

$$\ln(1+\lambda_i) = \frac{\sum_{a=25}^{99} \beta^{a-25}(S_i(a) - S_{US}(a))u_i(a)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)} + \frac{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)(u_i(a) - u_{US}(a))}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)}$$

Let $\Delta s_i(a) = \frac{\beta^{a-25} S_i(a) - \beta^{a-25} S_{US}(a)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)}$ and $s_{US}(a) = \frac{\beta^{a-25} S_{US}(a)}{\sum_{a=25}^{99} \beta^{a-25} S_{US}(a)}$. Then, we can write

$$\ln(1+\lambda_i) = \sum_a \Delta s_i(a) u_i(a) + \sum_a s_{US}(a)(u_i(a) - u_{US}(a))$$

Then write the equivalent variation as

$$\begin{aligned} \ln(1+\lambda_i) = & \sum_a \Delta s_i(a) u_i(a) \\ & + \sum_a s_{US}(a)(\ln c_i(a) - \ln c_{US}(a)) \\ & + \sum_a s_{US}(a)(v(l_i(a)) - v(l_{US}(a))). \end{aligned}$$

This is the same decomposition as in the text.

A.2 Deriving the Equivalent Variation when $\sigma \neq 1$

Suppose the utility function is

$$u(c,l) = \frac{\left[c \exp\left(\frac{\theta}{1+\frac{1}{\epsilon}}(1-l)^{1+\frac{1}{\epsilon}}\right)\right]^{1-\sigma} - 1}{1-\sigma}.$$

The equivalent variation is defined implicitly as

$$\sum_a \beta^{a-25} S_i(a) \left\{ \bar{u} + \frac{\left[c_i(a)(1+g)^{a-25} \exp\left(\frac{\theta}{1+\frac{1}{\epsilon}} (1-l_i(a))^{1+\frac{1}{\epsilon}} \right) \right]^{1-\sigma} - 1}{1-\sigma} \right\}$$

$$= \sum_a \beta^{a-25} S_{US}(a) \left\{ \bar{u} + \frac{\left[c_{US}(1+g)^{a-25}(a)(1+\lambda_i) \exp\left(\frac{\theta}{1+\frac{1}{\epsilon}} (1-l_{US}(a))^{1+\frac{1}{\epsilon}} \right) \right]^{1-\sigma} - 1}{1-\sigma} \right\}$$

Let $V_i = \sum_a \beta^{a-25} S_i(a) \left\{ \frac{\left[c_i(a)(1+g)^{a-25} \exp\left(\frac{\theta}{1+\frac{1}{\epsilon}} (1-l_i(a))^{1+\frac{1}{\epsilon}} \right) \right]^{1-\sigma}}{1-\sigma} \right\}$. Then, we have

$$V_i + \sum_a \beta^{a-25} S_i(a) \left[\bar{u} - \frac{1}{1-\sigma} \right] = (1+\lambda_i)^{1-\sigma} V_{US} + \sum_a \beta^{a-25} S_{US}(a) \left[\bar{u} - \frac{1}{1-\sigma} \right].$$

Thus,

$$(1+\lambda_i)^{1-\sigma} = \frac{V_i + \sum_a \beta^{a-25} \left[S_i(a) - S_{US}(a) \right] \left(\bar{u} - \frac{1}{1-\sigma} \right)}{V_{US}} \qquad \blacksquare$$

B The Geography of Economic Welfare

In a similar framework to ours, Falcettoni and Nygaard (2020) examine the distribution of economic welfare across states in the US. While we do not directly address the question of geography, we do have the demographic composition of each state from the ACS data. We estimate the welfare for a state by taking a weighted average of each demographic group's economic welfare with the weights being equal to the proportion of a demographic group in a state. We then compare our estimates for each state's economic welfare with the calculations in Falcettoni and Nygaard.

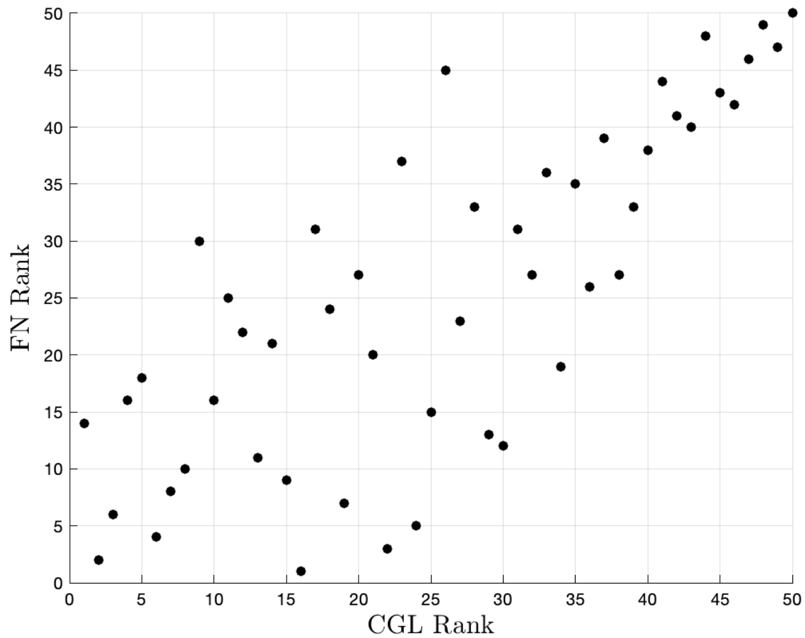

Figure B1: *Comparison of Two Welfare Measures*

Notes: Each states equivalent variation is $\lambda \times 100$ for each demographic group from Eq. (2) and weighted by each states demographic composition. The CGL rank is the states ranking in equivalent variation by highest to lowest and is plotted against the corresponding state rank in Falcettoni and Nygaard (2020).

The results are shown in Figure B1. The correlation between the two series is 0.78. The models are very closely aligned on the states with low economic welfare which is shown in the top right of the plot. Eight of the ten states with the lowest welfare in our model are also in the bottom 10 of Falcettoni and Nygaard's model, with Indiana and Georgia being the exceptions (38th and 40th respectively in Falcettoni and Nygaard). Our predictions are different at the top of the distribution where only five of the ten states with the highest welfare in our model are in Falcettoni and Nygaard's top ten.

This high correlation is somewhat surprising given that Falcettoni and Nygaard (2020) include the consumption of housing and also incorporate different prices for consumer goods across states while our measure is solely driven by demographic composition. Nevertheless, it is reassuring that we come up with a similar ranking.

C Data Description

C.1 Demographic Classification

For our data, we classify individuals according to sex, race/ethnicity, educational attainment, and age. In all, we have 16 independent groups (and ages within each group). Within our data, we split each group into males and females. We then group at the race/ethnicity level (from heron called *race*). These are non-Hispanic White (*White*), non-Hispanic Black (*Black*), non-Hispanic Asian (*Asian*), and Hispanics as their own category independent of race.[18] Within our data, we use "bridged-race" categories where an individual identifies as a single race, thereby avoiding complexities with mixed-race individuals. We note that Markides and Eschbach (2005) find evidence, using 1999 and 2000 data, that for Mexican immigrant men in particular may return to Mexico in old age, thereby increasing the longevity numbers in US statistics for this group.

Then within each sex and racial category, we then group by educational attainment. Our groups are less than a bachelor's degree and bachelor's degree or more. We opted to keep just two broad education categories to ensure our sample size is adequate. Moreover, we define age groups by five-year intervals from 25-24 to 95-99. Again, we do not use single age groups for sample size considerations. Table C1 lists the 16 independent groups (with ages 25-99 within each group).

Table C1: Individual Groups

Male			Female		
White	less than bachelor's	1	White	less than bachelor's	9
	bachelor's degree +	2		bachelor's degree +	10
Hispanic	less than bachelor's	3	Hispanic	less than bachelor's	11
	bachelor's degree +	4		bachelor's degree +	12
Black	less than bachelor's	5	Black	less than bachelor's	13
	bachelor's degree +	6		bachelor's degree +	14
Asian	less than bachelor's	7	Asian	less than bachelor's	15
	bachelor's degree +	8		bachelor's degree +	16

[18] We use Hispanics as a stand-in for those of Latin American descent (less Brazil). However, Spaniards and those of direct Spanish descent are included but we assume these constitute a small number of "Hispanics."

C.2 Survival probabilities

We calculate our survival probabilities from 2013-2017 by sex, age, race, and educational attainment. Our primary method follows Sasson (2016) but we update the coverage to more recent years and provide finer categories, specifically by race.[19] We use the US Multiple Cause of Death (MCD) micro-level mortality files in conjunction with the American Community Survey to capture a measure of age-specific survival probabilities.

American Community Survey

We use the American Community Survey (ACS) public use files to estimate the population by sex, age, race, and educational attainment for the population by each category. The ACS includes a "race-bridging" which classifies each person to a single race category. This is the approach in the mortality files which allows us to match race one-for-one with the MCD files. For 2015-2017, the ACS does not report single race categories so we follow the method by Liebler and Halpern-Manners (2008) to assign single race categories for each observation for these years.[20].

We begin by assigning each person to one of 15 age groups which are classified by five-year groups, ages 25-29, 30-34, ... , 90-94, 95-99.[21] Each person is then classified by age, sex, race, and education. The population counts in each category give us the person-years of exposure for each group.

Multiple Cause of Deaths files

We use the US Multiple Cause of Death (MCD) public use mortality files to estimate the number of deaths for each demographic category from 2013-2017. Sasson (2016) (and

[19]For example, Sasson (2016) notes that prior to 2000 (and to an extent 2010), accurate Hispanic counts in the death registry are not available. They are in our years.

[20]Program is found at https://usa.ipums.org/usa/resources/volii/race_bridge_stata_program.txt

[21]The age variable in the ACS is top-coded. See below how we address survival probabilities for ages 90-99.

other cited in his paper) note that there are some limitations to the data, particularly on educational attainment since this is reported by someone other than the deceased. Even so, it is still perhaps one of the only files that can be used to estimate mortality by fine demographic categories.

We begin again by assigning each person to one of 15 age groups which are classified by 5 year groups, ages 25-29, 30-34, ... , 90-94, 95-99. Each person is then classified by age, sex, race, and education. The population counts in each category give us the deaths for each group.

Although we classify education as a person's highest level or degree, prior to 2003 the MCD files report education by years. In 2003, states could report the deceased's education as either year based or highest level or degree. Not all states initially switched to the degree based classification of education, so the 2003 and 2004 data contains a mix of deaths reported by years and highest degree. To be consistent with our degree based classification of education in the paper, for those states reporting education using the year based system, we translate the years to education to whether they received a bachelor's or not.

The MCD does have a potentially significant amount of missing information by educational attainment. Therefore, we impute missing education values in the spirit of Sasson (2016). We assume all missing data on education is random.

Then, using the data on non-missing education values, we can estimate

$$P(Education|Death, X) \tag{8}$$

where X is a vector denoting age group, race, gender, and state. We find the conditional probabilities given in (8) for those who died in each year. We impute the missing education observations by randomly drawing from the conditional probabilities by X.

Once we estimate the survival probabilities using the ACS population and MCD death data, we note for the oldest age groups (ages 90-94 and 95-99) the ACS has top-coded ages, which give inaccurate survival probabilities in the latter two age groups. Within our model, we bound lifetimes to end at age 99, so in the data we impose survival probability to age 100 is zero and, using the age-profiles for survival we interpolate survival probabilities for ages 90-94 and 95-99 with cubic splines for each group classified on education, race, and gender.

C.3 Consumption Expenditure Survey Data

The data on consumption and leisure comes from the Consumer Expenditure Survey, compiled by the U.S. Bureau of Labor Statistics. The data is from the FMLI and MEMI files. The sample size for each demographic group is shown in Table C2.

Table C2: CE Unweighted Sample Size by Group, 2013-2017

		Sample Size	
Race/Ethnicity	Education	Male	Female
White	less than bachelor's	25,905	27,947
	bachelor's	22,645	24,361
Hispanic	less than bachelor's	8,237	8,965
	bachelor's	2,220	2,813
Black	less than bachelor's	4,403	5,688
	bachelor's	1,984	2,958
Asian	less than bachelor's	1,628	2,123
	bachelor's	2,718	2,939

Consumption expenditure

The CEX is a rotating survey where a household is interviewed up to five consecutive times. Respondents are interviewed once a quarter, but the interview can occur in any of the three months within that quarter.[22] Using household expenditures across interviews, we calculate total consumption for households that have data for all three months within a calendar quarter. We sum the relevant expenditures reported for the current quarter (CQ in CEX notation) and previous quarter (PQ) to get consumption by calendar quarter. We keep

[22]see, for example, http://data.nber.org/ces/2015/csxintvw.pdf, p.22.

only the observations that have consumption for the entire calendar quarter. The CEX Code in Table C3 shows the expenditure categories (for CQ – variable names for previous quarter analogously end in PQ).

We construct real non-durable consumption expenditures by deflating the household's expenditures for each of the categories listed in Table C3 by its category-specific deflator. The table also the BLS CPI code. Total household (consumer unit) consumption is the sum of the expenditures on components 1-13.

Table C3: CEX consumption categories and CPI categories

	CEX		CPI	
	Category Name	CEX Code (CQ)	Category Name	CPI Code
1	Food	FOODCQ	Food	SAF1
2	Alcohol beverages	ALCBEVCQ	Alcoholic beverages	SAF116
3	Tobacco	TOBACCCQ	Tobacco and smoking products	SEGA
4	Utilities	UTILCQ	Fuels and utilities	SAH2
5	Personal care	PERSCACQ	Personal care	SAG1
6	Household operations	HOUSOPCQ	Household furnishings and operations	SAH3
7	Public transportation	PUBTRACQ	Public transportation	SETG
8	Gas and motor oil	GASMOCQ	Motor fuels	SETB
9	Apparel	APPARCQ	Apparel	SAA
10	Miscellaneous expenditures	MISCCQ	Miscellaneous personal services	SEGD
11	Entertainment	ENTERTCQ	Entertainment	SAR
12	Rented dwelling	RENDWECQ	Rent of primary residence	SEHA
13	Owner occupied est. rent	RENTEQVX	Rent of primary residence	SEHA

Calculating Per-Person Consumption: Per person consumption is calculated by taking total household (consumer unit) consumption and adjusting by the OECD-modified equivalence scale[23] which assigns a value of one for the first member of the household, a value of 0.5 for each additional adult household member, and 0.3 for each child (we define as age 17 or below) by quarter. An observation is scaled by four to express consumption at an annual rate, consistent with the model. Further, we drop any observations that have zero food expenditures and have non-positive household income.

[23] see http://www.oecd.org/els/soc/OECD-Note-EquivalenceScales.pdf.

Income

We calculate per-person income by summing each individual's pre-tax wage and self-employment income and adding it to the household's income from other sources divided by the number of adults in the household in each time period

$$income_i = \frac{\text{individual income}_i \times 100}{CPI} + \frac{\frac{\text{household income}_j}{adults_j} \times 100}{CPI}$$

for individual i in household j. Table C4 lists the CEX components of individual and household income sources.

Table C4: CEX individual and household income

	CEX Code	Description
	Individual Income	
1	SALARYXM	Wage and salary income
2	NONFARMX	Income from nonfarm business
3	FARMINCX	Income from own farm
	Household Income	
1	FRRETIRX	Social security and railroad benefits
2	FSSIX	Supplemental security income
3	UNEMPLX	Unemployment compensation
4	COMPENSX	Workmen's compensation and veteran's payments
5	WELFAREX	Public assistance or welfare
6	INTEARNX	Interest on savings accounts or bonds
7	FININCX	Regular income from dividends, royalties, estates, or trusts
8	PENSIONX	Pensions or annuities from private companies
9	INCLOSSA	Income from roomers or boarders
10	INCLOSSB	Payments from rental properties
11	ALIOTHX	Alimony from persons outside consumer unit
12	CHDOTHX	Child support payments
13	OTHRINCX	Money for foster children, cash scholarships, fellowships, or stipends
14	JFDSTMPA	Value of food stamps

Top coded ages

The CEX data top codes ages. While it varies by year, it is in the mid 80s for our sample. For example, we have consumption for all individuals 85 and over, but cannot distinguish what 5 year age group the consumption is allocated to. Our method is to first predict consumption by age group for the older ages and then rescale these predicted values to be consistent with data. Specifically, we predict consumption by age group for age groups 12-15 and then scale this prediction so it matches the average aggregated consumption in the data for age groups 12-15 together. To allocate the consumption expenditures for the oldest age groups $a = \{12, ...15\}$, we first regress demographic characteristics for demographic groups $i = \{1, ...24\}$ on consumption by age a

$$
\begin{aligned}
\hat{c}_{i,a} &= \alpha + \beta_1 Sex_{i,a} * Age_{i,a} + \beta_2 Sex_{i,a} * Age_{i,a}^2 + \beta_3 Educ_{i,a} * Age_{i,a} + \beta_4 Educ_{i,a} * Age_{i,a}^2 \\
&+ \beta_5 Race_{i,a} * Age_{i,a} + \beta_6 Race_{i,a} * Age_{i,a}^2 + \zeta_{i,a}
\end{aligned}
$$

Since the data for the older ages is aggregated, we know average consumption for the older ages in age groups 12-15. For demographic group i, call this $\bar{c}_i$. Next, we can calculate the average predicted consumption for ages 12-15. For demographic group i, call this $\bar{\hat{c}}_i$, which is the weighted average of consumption for age groups 12-15. The weights are the population weights from the ACS, under the assumption that the CEX data is representative of US population shares. This is calculated as $\bar{\hat{c}}_i = \sum_{a=12}^{15} \hat{c}_{i,a} p_{i,a}$, where $p_{i,a}$ is the population share of age group a and demographic group i. We can then use this as a scaling factor to see how close the predicted consumption by age group is to the data, $\frac{\bar{c}_i}{\bar{\hat{c}}_i}$. We use this scaling factor the then rescale consumption by single age group so the average of age groups 12-15 equals that in the data. That is, consumption by age group is given by $\frac{\bar{c}_i}{\bar{\hat{c}}_i} \hat{c}_{i,a}$.

C.4 Leisure

We use the American Time Use Survey for data on market work, child care, and home production. We download the microdata from the IPUMS American Time Use Survey extract builder (Sandra L. Hofferth and Backman, 2020).

We follow Jones and Klenow (2016) and let there by 5840 hours available for working a year, which is 16 hours per day for 365 days a year. Let annual hours in market production, home production, and childcare be denoted as h_{work}, h_{home}, and h_{child}, respectively. Leisure is then $1 - \frac{h_{work} + h_{home} + h_{child}}{365 \times 16}$. We further impose that individuals age 85 spend no hours in market production (i.e. $h_{work} = 0$).

The ATUS uses a six digit coding structure to classify all time use activities. The first two digits are the broadest, the second two a little more specific, and the final two even more specific. For instance, 030101 is classified under "caring for and helping household members" (the 03), and further classified under "caring for and helping children" (0301), and finally defined as "physical care for household children" (030101).

We include the following activities in market work:

1. Working (050100).
2. Work-related activities (050200).
3. Travel related to work (180501).
4. Travel related to work-related activities (180502).

We include the following activities in home production:

1. All household activities (020000) except for those related to animals and pets (020600).
2. Caring for (03400) and helping (03500) household adults.
3. All activities related to consumer purchases (070000).
4. All activities related to professional and personal care services (080000) except for medical care services (080400).
5. All household services (090000).
6. All travel times associated with the preceding activities (18XX00).

We include the following activities in child care:

1. Caring for and helping household children (030100).
2. Activities related to household children's education (030200).

3. Activities related to household children's health (030300).

4. All travel times associated with the preceding activities (18XX00).